

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims

1. (Currently amended) A method for processing network discovery data, comprising the following operations:

defining a plurality of network data aggregations, wherein the network data aggregations are based on switch boundaries, wherein interswitch links are used to define the plurality of network data aggregations, and wherein the network data aggregations are defined via an XML file using a DTD format, wherein the network data aggregations are based on zoning information including a zone that is a set of ports in a fabric that make up a logical boundary, wherein an inband agent executing in a storage management server conducts discovery by issuing queries to switches in a storage area network via an Internet Protocol network segment, and wherein an outband agent executing in a network management server conducts discovery by issuing queries via communication pathways in the storage area network;

computing a current state value for at least one of the network data aggregations, wherein the current state value is a CRC code, and wherein the CRC code is computed utilizing data associated with the corresponding network data aggregation and a CRC polynomial;

for at least one current state value, determining if the current state value is different than a corresponding prior state value for a corresponding network data aggregation, wherein the operation of, for at least one current state value, determining if the current state value is different than a corresponding prior state value, is performed for each of a plurality of levels in a hierarchy of data aggregations, wherein the plurality of data aggregations includes at least one data aggregation that is a subset of a corresponding superset data aggregation, and wherein the subset data aggregation is located in a hierarchal ordering after the corresponding superset data aggregation; and

merging data corresponding with at least one network data aggregation determined to have a current state value that is different than a corresponding prior state value, with prior data corresponding with at least one different network data aggregation determined to

have a current state value that is not different than a corresponding prior state value for the different network data aggregation.

2-7. (Canceled)

8. (Previously presented) The method of claim 1, wherein the operation of computing a current state value for at least one of the data aggregations is performed by at least one agent discovery service.

9. (Canceled)

10. (Previously presented) The method of claim 1, wherein the operation of computing a current state value for at least one of the data aggregations comprises processing data in the at least one of the data aggregations in a prescribed order.

11. (Previously presented) The method of claim 1, wherein the operations further comprise, prior to the operation of computing a current state value for at least one of the data aggregations, organizing data in the at least one of the data aggregations in a prescribed order.

12-17. (Canceled)

18. (Previously presented) The method of claim 1, wherein the operations further comprise requesting polling on data aggregations that are subsets of a superset data aggregation that has a changed state value.

19-36. (Canceled)

37. (Previously presented) The method of claim 1, wherein the switch boundaries are defined at links that are chosen to minimize a number of links at the switch boundaries.

38. (Currently amended) A system, comprising:

a memory; and

a processing device coupled to the memory, wherein the processing device is programmed to perform operations for processing network discovery data, the operations comprising:

defining a plurality of network data aggregations, wherein the network data aggregations are based on switch boundaries, wherein interswitch links are used to define the plurality of network data aggregations, and wherein the network data aggregations are defined via an XML file using a DTD format, wherein the network data aggregations are based on zoning information including a zone that is a set of ports in a fabric that make up a logical boundary, wherein an inband agent executing in a storage management server conducts discovery by issuing queries to switches in a storage area network via an Internet Protocol network segment, and wherein an outband agent executing in a network management server conducts discovery by issuing queries via communication pathways in the storage area network;

computing a current state value for at least one of the network data aggregations, wherein the current state value is a CRC code, and wherein the CRC code is computed utilizing data associated with the corresponding network data aggregation and a CRC polynomial;

for at least one current state value, determining if the current state value is different than a corresponding prior state value for a corresponding network data aggregation, wherein the operation of, for at least one current state value, determining if the current state value is different than a corresponding prior state value, is performed for each of a plurality of levels in a hierarchy of data aggregations, wherein the plurality of data aggregations includes at least one data aggregation that is a subset of a corresponding superset data aggregation, and wherein the subset data aggregation is located in a hierarchal ordering after the corresponding superset data aggregation; and

merging data corresponding with at least one network data aggregation determined to have a current state value that is different than a corresponding prior

state value, with prior data corresponding with at least one different network data aggregation determined to have a current state value that is not different than a corresponding prior state value for the different network data aggregation.

39. (Previously presented) The system of claim 38, wherein the operation of computing a current state value for at least one of the data aggregations is performed by at least one agent discovery service.

40. (Previously presented) The system of claim 38, wherein the operation of computing a current state value for at least one of the data aggregations comprises processing data in the at least one of the data aggregations in a prescribed order.

41. (Previously presented) The system of claim 38, wherein the operations further comprise, prior to the operation of computing a current state value for at least one of the data aggregations, organizing data in the at least one of the data aggregations in a prescribed order.

42-43. (Canceled).

44. (Previously presented) The system of claim 38, wherein the operations further comprise requesting polling on data aggregations that are subsets of a superset data aggregation that has a changed state value.

45. (Previously presented) The system of claim 38, wherein the switch boundaries are defined at links that are chosen to minimize a number of links at the switch boundaries.

46. (Currently amended) A computer readable storage medium tangibly embodying a program of computer-readable instructions executable by a digital processing apparatus to perform operations for processing network discovery data, the operations comprising:

defining a plurality of network data aggregations, wherein the network data aggregations are based on switch boundaries, wherein interswitch links are used to define the plurality of

network data aggregations, and wherein the network data aggregations are defined via an XML file using a DTD format, wherein the network data aggregations are based on zoning information including a zone that is a set of ports in a fabric that make up a logical boundary, wherein an inband agent executing in a storage management server conducts discovery by issuing queries to switches in a storage area network via an Internet Protocol network segment, and wherein an outband agent executing in a network management server conducts discovery by issuing queries via communication pathways in the storage area network;

computing a current state value for at least one of the network data aggregations, wherein the current state value is a CRC code, and wherein the CRC code is computed utilizing data associated with the corresponding network data aggregation and a CRC polynomial;

for at least one current state value, determining if the current state value is different than a corresponding prior state value for a corresponding network data aggregation, wherein the operation of, for at least one current state value, determining if the current state value is different than a corresponding prior state value, is performed for each of a plurality of levels in a hierarchy of data aggregations, wherein the plurality of data aggregations includes at least one data aggregation that is a subset of a corresponding superset data aggregation, and wherein the subset data aggregation is located in a hierarchal ordering after the corresponding superset data aggregation; and

merging data corresponding with at least one network data aggregation determined to have a current state value that is different than a corresponding prior state value, with prior data corresponding with at least one different network data aggregation determined to have a current state value that is not different than a corresponding prior state value for the different network data aggregation.

47. (Previously presented) The computer readable storage medium of claim 46, wherein the operation of computing a current state value for at least one of the data aggregations is performed by at least one agent discovery service.

48. (Previously presented) The computer readable storage medium of claim 46, wherein the operation of computing a current state value for at least one of the data aggregations comprises processing data in the at least one of the data aggregations in a prescribed order.

49. (Previously presented) The computer readable storage medium of claim 46, wherein the operations further comprise, prior to the operation of computing a current state value for at least one of the data aggregations, organizing data in the at least one of the data aggregations in a prescribed order.

50-51. (Canceled)

52. (Previously presented) The computer readable storage medium of claim 46, wherein the operations further comprise requesting polling on data aggregations that are subsets of a superset data aggregation that has a changed state value.

53. (Previously presented) The computer readable storage medium of claim 46, wherein the switch boundaries are defined at links that are chosen to minimize a number of links at the switch boundaries.